

The Stretch - Length Tradeoff in Geometric Networks: Average Case and Worst Case Study

David Aldous
Department of Statistics
U.C. Berkeley CA 94720

Tamar Lando
Department of Philosophy
Columbia University
1150 Amsterdam Ave
New York NY 10027

April 11, 2014

Abstract

Consider a network linking the points of a rate-1 Poisson point process on the plane. Write $\Psi^{\text{ave}}(s)$ for the minimum possible mean length per unit area of such a network, subject to the constraint that the route-length between every pair of points is at most s times the Euclidean distance. We give upper and lower bounds on the function $\Psi^{\text{ave}}(s)$, and on the analogous “worst-case” function $\Psi^{\text{worst}}(s)$ where the point configuration is arbitrary subject to average density one per unit area. Our bounds are numerically crude, but raise the question of whether there is an exponent α such that each function has $\Psi(s) \asymp (s-1)^{-\alpha}$ as $s \downarrow 1$.

1 Introduction

The topic *geometric spanner networks* [13] concerns design of networks on arbitrary sets of vertices in the plane (or higher dimensions). The interpretation of “size” of the network is sometimes as number of edges and sometimes as network length (sum of Euclidean edge lengths). Similarly, the interpretation of within-network distance between two vertices v, w is sometimes taken as minimum number of edges of a route between them (*hop length*) and sometimes as shortest total length (sum of Euclidean edge lengths) of a route between them (*route length* $r(v, w)$). In the latter setting, how well the network provides short routes is often measured by a statistic such as

$$S := \max_{v \neq w} \frac{r(v, w)}{d(v, w)} \geq 1 \quad (1)$$

where $d(v, w)$ denotes straight line (Euclidean) distance. The statistic S is called the *stretch* or *spanning ratio* of the network, and a network with stretch S is called an *S-spanner*.




Figure 1: Illustration of possible networks. The left diagram shows a network on 4 cities which is connected (a driver can switch roads at the junction where they cross), the center diagram shows that other junctions may be created, and the right diagram (envisaged as part of a larger network) shows that roads need not be closely related to cities at all.

Most work on this topic has emphasized algorithms – either algorithms for constructing spanners, or the use of spanners in algorithms for computational geometry problems. We address a more fundamental geometric question: what is the tradeoff between stretch and network length? In formulating a mathematical question we have in mind the example of an inter-city road network (rather than, say, a wireless communication network) and indeed we find it helpful to use the vivid natural language of *cities*, *roads*, *junctions* in place of the mathematical language of *vertices*, *edges*, *Steiner points*.

Our notion of “network” is the general notion suggested by real-world road networks, illustrated in Fig. 1. The only implicit convention is that networks be connected and consist of line segments (rather than curves). Note this convention differs from the more familiar and more restrictive assumption that an edge can only be the line segment (v, w) between two of the given vertices. There does not seem to be standard terminology to emphasize the distinction: we will write *Steiner network* for our setting and *graph network* for the more restrictive assumption.

Our underlying setting is a configuration of n cities at arbitrary positions $\mathbf{z}_n = (z_1, \dots, z_n)$ in a square of area n . For a network \mathcal{N} connecting these cities, write $S(\mathcal{N})$ for the statistic (1) and write

$$L(\mathcal{N}) = \frac{1}{n} \times (\text{network length of } \mathcal{N})$$

for normalized network length. We then define

$$\psi_n(\mathbf{z}_n, s) := \inf\{L(\mathcal{N}) : S(\mathcal{N}) \leq s\}$$

the infimum over all networks \mathcal{N} connecting the cities \mathbf{z}_n . So this quantifies the optimal trade-off between length and stretch for the given configuration. We can now consider in parallel the worst-case, that is $\sup_{\mathbf{z}_n} \psi_n(\mathbf{z}_n, s)$, and the average case $\mathbb{E}\psi_n(\mathbf{Z}_n, s)$, where \mathbf{Z}_n consists of n independent uniform random

positions in the area- n square. The purpose of this set-up is that it is intuitively obvious that there must exist limit functions

$$\Psi^{\text{worst}}(s) = \lim_{n \rightarrow \infty} \sup_{\mathbf{z}_n} \psi_n(\mathbf{z}_n, s)$$

$$\Psi^{\text{ave}}(s) = \lim_{n \rightarrow \infty} \mathbb{E} \psi_n(\mathbf{Z}_n, s)$$

where $0 < \Psi^{\text{ave}}(s) \leq \Psi^{\text{worst}}(s) \leq \infty$ for $1 < s < \infty$.

The goal of this paper is to study the functions Ψ^{ave} and Ψ^{worst} . As probabilists the authors are primarily interested in the average-case setting, but it seems natural to treat the worst-case setting in parallel. Here is what we shall do.

- Prove existence of these limit functions (section 2).
- Prove that their $s \rightarrow \infty$ limits are equal to (rather than greater than) the associated Steiner tree constants (section 3).
- See what upper bounds can be derived from known results for graph networks (sections 1.1 and 4).
- Derive upper bounds on $\Psi^{\text{worst}}(s)$ from elementary constructions where one first lays down a regular network of roads without paying attention to city positions, and then adds local links from cities to the network (section 5).
- Derive upper bounds on $\Psi^{\text{ave}}(s)$ from Steiner network analogs of the θ -graphs discussed in section 4 (section 6).
- Derive lower bounds on $\Psi^{\text{ave}}(s)$ for small s , based on the stochastic geometry relationship between network length and rate of intersections with a typical line (section 7).
- Derive lower bounds on $\Psi^{\text{worst}}(s)$ based on a notion of “local optimality” for specific networks on specific configurations (section 8).

The sections are, to a large extent, independent of each other. Our description of $\Psi^{\text{ave}}(s)$ as an $n \rightarrow \infty$ limit was intended to facilitate comparison with the worst-case setting. A more abstract interpretation of $\Psi^{\text{ave}}(s)$ in terms of networks on a Poisson point process on \mathbb{R}^2 is given in section 2.1.

Getting explicit values for these functions analytically seems impossible. Even getting convincing numerical values would provide a challenge for designers of heuristic algorithms, and we have not attempted to do so. Our bounds are numerically crude; this paper is intended to initiate study of these functions, not to give definitive results.

An interesting theoretical question that seems more amenable to analytic study is the scaling behavior in the $s \downarrow 1$ limit. That is, in the spirit of “universality” in statistical physics, one can speculate that there exists an exponent α such that

$$\Psi(s) \asymp (s - 1)^{-\alpha} \text{ as } s \downarrow 1$$

where the value of α does not depend on any detailed assumptions in the model (worst-case or average-case; the Steiner network case or the graph network case) but instead depends only on the fact that we are studying the length-stretch trade-off in two-dimensional space. Our results imply crude bounds on α : an upper bound of $\frac{3}{4}$ for Ψ^{ave} (Corollary 6.5) and $\frac{5}{4}$ for Ψ^{worst} (24), and a lower bound of $\frac{3}{8}$ for Ψ^{ave} and hence for Ψ^{worst} also (Proposition 7.1). But considering these $s \rightarrow 0$ limits is considering increasingly dense networks (“covering the countryside in tarmac”, we say in talks), which is hard to motivate.

1.1 What is already known?

The literature on geometric spanner networks focusses on worst-case bounds on stretch (and many other statistics of networks) produced by algorithmic procedures from arbitrary configurations, works in the “graph network” setting rather than our Steiner network setting, and emphasizes hop length more than route length. So while the techniques of that field are clearly relevant, it is not so easy to directly apply their results to the study of Ψ^{ave} and Ψ^{worst} .

One relevant result in [13] is Theorem 15.2.16, which says that for small $s > 0$ one can construct $(1 + s)$ -spanners such that (amongst other properties) the network length is bounded by $O(s^{-4} \times (\text{length of MST}))$, where MST denotes the minimum spanning tree on the given configuration. It is well-known and elementary (see e.g. [15] section 2.2) that in the worst case the length of MST is $O(n)$, so the theorem mentioned above implies

$$\Psi^{\text{worst}}(s) = O((s - 1)^{-4}) \text{ as } s \downarrow 1 \quad (2)$$

and in particular that $\Psi^{\text{worst}}(s) < \infty$ for all $s > 1$.

A remarkable result [9] is that the Delaunay triangulation is always a t -spanner for $t = \frac{2\pi}{3 \cos \pi/6} \approx 2.42$. Because the length of a Delaunay triangulation is *not* $O(n)$ in the worst case, this does not directly help us bound Ψ^{worst} . But in the random model, a classical result ([11] page 113) shows the limit normalized length of the Delaunay triangulation equals $\frac{32}{3\pi} = 3.40\dots$. So we get a numerical bound

$$\Psi^{\text{ave}}(2.42) \leq 3.40. \quad (3)$$

More generally, there are known bounds [6] on stretch for the well-studied one-parameter θ -graph family of graph networks. As above, their lengths are not $O(n)$ in the worst case, so this does not directly help us bound Ψ^{worst} . But again we can calculate mean lengths in the random model, and so deduce explicit upper bounds on $\Psi^{\text{ave}}(s_m)$ for a certain sequence $s_m \downarrow 1$ (section 4). However for small s we get better bounds, in our Steiner network setting, from the construction in section 6.

1.2 Other statistics for route-length efficiency

In defining a statistic to summarize the effectiveness of a network in providing short routes, one may be more interested in the typical value of

$$R(v, w) = \frac{r(v, w)}{d(v, w)} - 1$$

than in the maximum value used in the definition of *stretch*. One might first consider the summary statistic $\text{ave}_{v,w} R(v, w)$, which somewhat counter-intuitively can easily be made very small for large n [2, 8]. It is argued in [4] that the most appropriate summary statistic R is defined as follows. For each distance d , set $\rho(d) = \text{average of } R(v, w) \text{ over city-pairs at distance approximately } d$; then let R be the maximum of $\rho(d)$ as d varies. The trade-off between R and normalized length, in the average-case setting, is discussed in [4], and the motivation for the present paper was to make a connection with the topic of spanner networks.

2 Existence of the limit functions Ψ

In this section we use a subadditivity argument in the spirit of [15, 18] to prove existence of the limit functions Ψ . Note that in the most familiar kind of spatial subadditivity argument a big square is divided into small subsquares, and optimal solutions on subsquares are used to construct some near-optimal solution on the big square. We argue in the opposite direction: use an optimal solution on the big square to construct near-optimal solutions on subsquares. This leads to the “superadditive” inequalities (5, 7).

Fix $1 < s < \infty$ and let a_n be the worst-case value (over configurations $\mathbf{z} = (z_1, \dots, z_n)$ of cities in the square of area n) of the length of the shortest network on \mathbf{z} with stretch $\leq s$. We shall prove existence of the limit

$$\Psi^{\text{worst}}(s) := \lim_n n^{-1} a_n \leq \infty. \quad (4)$$

We will first argue

$$\frac{a_n}{n} \leq \frac{a_{nk^2}}{nk^2} + \frac{4}{\sqrt{n}}, \quad n \geq 1, k \geq 2. \quad (5)$$

Fix n and k . Let \mathbf{z} be a configuration in the area- n square attaining a_n . Take k^2 copies of this configuration, and translate each to construct a configuration \mathbf{z}^* of nk^2 cities in the square of area nk^2 . By definition of a_{nk^2} there is a network on \mathbf{z}^* with stretch $\leq s$ and with length $\leq a_{nk^2}$. Add to this network the four boundary edges of each of the k^2 subsquares (so we get two copies of each edge interior to the big square). We now have a network \mathcal{N}^* whose length $c_{n,k}$ satisfies $c_{n,k} \leq a_{nk^2} + 4n^{1/2}k^2$. Consider the restriction of this network to one of the subsquares. The length of the restricted network may depend on the subsquare, but there must be at least one subsquare Q for which this length of the restricted network \mathcal{N}_Q is at most the average $c_{n,k}/k^2$. Routes in the network \mathcal{N}^* between cities of Q might go outside Q , but replacing these

external segments by the boundary edges of Q can only shorten the route length, so \mathcal{N}_Q has stretch at most s . But \mathcal{N}_Q defines (by translation) a network on the original configuration \mathbf{z} , and so $a_n \leq c_{n,k}/k^2$, which gives (5).

Deducing existence of a limit from (5) is one of many variants of routine “subadditivity” arguments, as follows. First note that (a_n) is increasing; indeed

$$a_{n+1} \geq a_n \sqrt{\frac{n+1}{n}}$$

by adding an arbitrary city to the configuration attaining a_n and rescaling. (Note this is one of many minor ways in which Steiner networks are technically more tractable than graph networks). Next define

$$\gamma = \liminf_n a_n/n \leq \infty$$

and use monotonicity to show

$$\gamma = \liminf_k \frac{a_{nk^2}}{nk^2}, \quad \text{for each fixed } n.$$

Then (5) shows

$$a_n/n \leq \gamma + 4n^{-1/2}$$

and so $\limsup_n a_n/n \leq \gamma$, meaning that indeed $\lim_n a_n/n = \gamma$.

This argument shows $\Psi^{\text{worst}}(s) < \infty$ exists. As mentioned at (2), existing results then imply $\Psi^{\text{worst}}(s) < \infty$ for all $s > 1$; this alternatively could be derived from the more elementary constructions in our section 5.

For the random model we use the same construction with a Poissonized number of random points. Fix s again. Let b_n be the expectation of the length of the shortest network with stretch $\leq s$ over n uniform random cities in the *unit* square. So $n^{1/2}b_n$ is the corresponding expectation in the area- n square. We shall prove existence of the limit

$$\Psi^{\text{ave}}(s) := \lim_n \frac{n^{1/2}b_n}{n}. \quad (6)$$

Write $N(t)$ for a random variable with Poisson(t) distribution and write

$$\beta_t = \mathbb{E}b_{N(t)}.$$

Take a Poisson point process (rate 1 per unit area) of cities on the whole plane. Now $t^{1/2}\beta_t$ is the expectation of the length of the shortest network with stretch $\leq s$ on the Poisson cities in an area- t square. Consider partitioning a square of area tk^2 into k^2 subsquares of area t . Repeating the argument for (5), now using a random subsquare, gives an inequality analogous to (5):

$$\frac{\beta_t}{t^{1/2}} \leq \frac{\beta_{tk^2}}{t^{1/2}k} + \frac{4}{t^{1/2}}, \quad 0 < t < \infty, \quad k \geq 2. \quad (7)$$

Using the fact that β_t is increasing in t , we can repeat the “subadditivity” argument to show existence of the limit

$$\lim_t t^{-1/2} \beta_t = \gamma^* \leq \infty. \quad (8)$$

The “average case better than worst case” inequality $n^{1/2} b_n \leq a_n$ and the fact $\Psi^{\text{worst}}(s) < \infty$ now easily imply $\gamma^* < \infty$.

To finish we need a routine “dePoissonization” argument to show that (7) and monotonicity of b_n imply

$$\lim_n n^{-1/2} b_n = \gamma^*. \quad (9)$$

First fix $\varepsilon > 0$ and consider $t_n/n \rightarrow 1 + \varepsilon$. Then

$$\beta_{t_n} \geq b_n \mathbb{P}(N(t_n) \geq n) = b_n(1 - o(1))$$

so

$$\limsup_n n^{-1/2} b_n \leq \limsup_n n^{-1/2} \beta_{t_n} = (1 + \varepsilon)^{1/2} \gamma^*$$

and the upper bound for (9) follows. Next, the following property of the Poisson distribution

$$\max_{i \geq n} \frac{\mathbb{P}(N((1 - \varepsilon)n) = i)}{\mathbb{P}(N(n) = i)} \rightarrow 0$$

implies

$$\mathbb{E} b_{N((1 - \varepsilon)n)} \mathbb{1}(N((1 - \varepsilon)n) \geq n) = o(\mathbb{E} b_{N(n)}) = o(n^{1/2})$$

and so

$$\beta_{(1 - \varepsilon)n} \leq b_n + o(n^{1/2})$$

implying

$$\liminf_n n^{-1/2} b_n \geq \liminf_n n^{-1/2} \beta_{(1 - \varepsilon)n} = (1 - \varepsilon)^{1/2} \gamma^*$$

and the lower bound for (9) follows.

2.1 Poisson process interpretation of Ψ^{ave}

The argument above interprets $\Psi^{\text{ave}}(s)$ as an $n \rightarrow \infty$ limit of the random n -city model. By standard weak convergence arguments which we will not give here (see e.g. [3] section 3.5 for more details in a somewhat similar setting) we can give an “exact” interpretation of $\Psi^{\text{ave}}(s)$ in terms of a Poisson (rate 1) point process of cities on the infinite two-dimensional plane. Consider a network \mathcal{N}_∞ on such cities whose distribution μ is translation invariant and ergodic. Associated with μ are two numbers: the stretch, say $S(\mu)$, and the normalized length (mean length-per-unit area), say $L(\mu)$, which is well-defined by translation invariance (of course these numbers might be $+\infty$). Then

$$\Psi^{\text{ave}}(s) = \inf\{L(\mu); \mu \text{ is translation invariant, } S(\mu) \leq s\}. \quad (10)$$

3 Short networks and the Steiner constants

Write $\mathbf{z}_n = (z_1, \dots, z_n)$ for a configuration of city positions in the square of area n . Write $\text{ST}(\mathbf{z}_n)$ for the Steiner tree (i.e. minimum length connected network) on \mathbf{z}_n , and for any network \mathcal{N} write $\text{len}(\mathcal{N})$ for its total length. By an easy “superadditive” argument similar to that in section 2, there exists a limit constant for worst-case normalized Steiner tree length:

$$c_{\text{worst}} := \limsup_{n \rightarrow \infty} \sup_{\mathbf{z}_n} n^{-1} \text{len}(\text{ST}(\mathbf{z}_n)). \quad (11)$$

It is known [7] that $c_{\text{worst}} \leq 0.995$ and that (by considering the hexagonal lattice) $c_{\text{worst}} \geq (3/4)^{1/4} = 0.9306$. Clearly we must have $\Psi^{\text{worst}}(s) \geq c_{\text{worst}}$ for all s , and this inequality must persist in the limit (which exists by monotonicity): $\lim_{s \rightarrow \infty} \Psi^{\text{worst}}(s) \geq c_{\text{worst}}$.

Turning to the average-case setting, it follows from the general theory of subadditive Euclidean functionals [15, 18] that there exists a limit constant c_{ave} such that

$$n^{-1} \text{len}(\text{ST}(Z_1, Z_2, \dots, Z_n)) \rightarrow c_{\text{ave}} \text{ in } L^1 \quad (12)$$

where the (Z_i) are independent uniform random in the area- n square. As above, we clearly have $\lim_{s \rightarrow \infty} \Psi^{\text{ave}}(s) \geq c_{\text{ave}}$. It is natural to guess (but not obvious) that these limit inequalities are really equalities. This guess is correct, as an immediate corollary of the following estimate for arbitrary city configurations.

Proposition 3.1. *There exists a function $\delta(s) \leq \infty$ with $\lim_{s \rightarrow \infty} \delta(s) = 0$, and a function $K(s) < \infty$, such that for all $1 < s < \infty$, all $n \geq K(s)$ and all city configurations \mathbf{z}_n in the area- n square, there exists a network \mathcal{N} connecting cities \mathbf{z}_n such that*

$$\text{stretch}(\mathcal{N}) \leq s; \quad n^{-1}(\text{len}(\mathcal{N}) - \text{len}(\text{ST}(\mathbf{z}_n))) \leq \delta(s).$$

Corollary 3.2.

$$\lim_{s \rightarrow \infty} \Psi^{\text{worst}}(s) = c_{\text{worst}} \text{ and } \lim_{s \rightarrow \infty} \Psi^{\text{ave}}(s) = c_{\text{ave}}.$$

The idea of the proof is to partition the area- n square into rectangles containing at most K cities, and then use a crude construction (Lemma 3.3) of networks on K cities. We will set up some notation, state the lemma, give the reduction of the Proposition to the lemma, and then prove the lemma.

Fix $K \geq 0$. Let A be a rectangle; write ∂A for its boundary, so that $\text{len}(\partial A)$ is its boundary length. Let y_1, \dots, y_K be an arbitrary configuration of K cities in A . Consider a network $\mathcal{N} = \mathcal{N}(A)$ in A which includes the boundary ∂A and links the cities to the boundary. For such a network define

$$\text{stretch}^*(\mathcal{N}) = \max_{y \neq y'} \frac{\text{route-length from } y \text{ to } y'}{d(y, y')} \quad (13)$$

where y and y' run over the cities and over points of ∂A .

Lemma 3.3. *Let $\hat{\mathbf{t}}$ be the Steiner tree on the cities y_1, \dots, y_K in a rectangle A and (possibly) other cities outside A . Let \mathbf{t} be the intersection of $\hat{\mathbf{t}}$ with A . There there exists a network \mathcal{N} in A containing \mathbf{t} and ∂A and linking the cities to ∂A , such that*

$$\text{stretch}^*(\mathcal{N}) \leq \rho(K); \quad \text{len}(\mathcal{N}) - \text{len}(\mathbf{t}) \leq 2 \text{len}(\partial A)$$

where $\rho(K) < \infty$ depends only on $K \geq 0$.

Proof of Proposition 3.1. Fix K and $n > K$. We use a simple decomposition, the *multidimensional search tree* or *k-d tree* [14]. Split the square $[0, n^{1/2}]^2$ into two rectangles using a vertical line through the city with median x -coordinate (if n is odd) or a vertical line separating the two median x -coordinate cities (if n is even). In either case, each rectangle has at most $n/2$ cities in its interior. Separately for each rectangle, split it into two rectangles using horizontal lines through the median y -coordinate(s). Now (end of stage 1) we have 4 rectangles, each with at most $n/4$ cities in its interior. Continue recursively for L stages, where L is the smallest integer such that $n4^{-L} \leq K$, to get a partition into 4^L rectangles, each with at most K cities in its interior. Write A for a generic rectangle in this partition.

Given a configuration \mathbf{z}_n in $[0, n^{1/2}]^2$, apply Lemma 3.3 (where $\hat{\mathbf{t}}$ is the Steiner tree on \mathbf{z}_n) to each A and the cities inside A to obtain a network $\mathcal{N}(A)$ satisfying

$$\text{stretch}^*(\mathcal{N}(A)) \leq \rho(K); \quad \text{len}(\mathcal{N}(A)) - \text{len}(\text{ST}(\mathbf{z}_n) \cap A) \leq 2 \text{len}(\partial A). \quad (14)$$

Then consider the network \mathcal{N} on the cities \mathbf{z}_n obtained as the union of networks $\mathcal{N}(A)$. Note that the bound $\rho(K)$ on $\text{stretch}^*(\mathcal{N}(A))$ does not depend on A . For any pair of cities z_i, z_j , we can define a route in \mathcal{N} between them by considering the points v_1, v_2, v_3, \dots at which a straight line between them intersects boundaries of successive rectangles A_1, A_2, A_3, \dots , and within each such rectangle A use the shortest route in $\mathcal{N}(A)$ between these boundary points (or the cities z_i, z_j themselves, at the ends). It follows that $\text{stretch}(\mathcal{N}) \leq \max_A \text{stretch}^*(\mathcal{N}(A)) \leq \rho(K)$. Note that the intermediate rectangles may contain no cities of \mathbf{z}_n , explaining why we must allow $K = 0$ in Lemma 3.3.

As a preliminary to bounding $\text{len}(\mathcal{N})$, we need to consider the total length of lines used in the original decomposition. Include a stage 0 in which the edges of the external boundary ∂_0 of $[0, n^{1/2}]^2$ are added. At stage 1, the length of lines added equals $2n^{1/2}$, and inductively at stage j the length of lines added equals $2^j n^{1/2}$. Because each segment of these added lines (except the external boundary) is part of the boundary of exactly 2 of the final rectangles A ,

$$\sum_A \text{len}(\partial A) = \text{len}(\partial_0) + 2 \sum_{j=1}^L 2^j n^{1/2} = 2n^{1/2} \left(2 + \sum_{j=1}^L 2^j \right) = 2^{L+2} n^{1/2}.$$

By definition of L we have $n4^{-(L-1)} > K$, giving $2^L \leq 2n^{1/2}K^{-1/2}$, and so

$$\sum_A \text{len}(\partial A) \leq 8nK^{-1/2}. \quad (15)$$

So

$$\begin{aligned} \text{len}(\mathcal{N}) &\leq \sum_A \text{len}(\mathcal{N}(A)) \\ &\leq \sum_A (\text{len}(\text{ST}(\mathbf{z}_n) \cap A) + 2\text{len}(\partial A)) \text{ by (14)} \\ &= \text{len}(\text{ST}(\mathbf{z}_n)) + 2 \sum_A \text{len}(\partial A). \end{aligned}$$

Combining with (15),

$$n^{-1}(\text{len}(\mathcal{N}) - \text{len}(\text{ST}(\mathbf{z}_n))) \leq 16K^{-1/2}.$$

We may assume $\rho(K) \uparrow \infty$ as $K \uparrow \infty$, and now Proposition 3.1 holds for $K(s) := \max\{K : \rho(K) \leq s\}$ and $\delta(s) := 16K^{-1/2}(s)$.

Proof of Lemma 3.3. We may suppose A is an $a_1 \times a_2$ rectangle, where $a_1 \leq a_2$. The network \mathcal{N} will consist of

- (i) \mathbf{t} (the intersection of $\hat{\mathbf{t}}$ with A)
- (ii) the boundary ∂A of A
- (iii) extra edges, of total length at most $\text{len}(\partial A)$.

Set $m = \lfloor a_2/a_1 \rfloor$ and partition A into $m+1$ similar $a_1 \times \frac{a_2}{m+1}$ rectangles by using m equally spaced roads of length a_1 . So the total length of these added roads is $ma_1 \leq a_2 \leq \frac{1}{2}\text{len}(\partial A)$. So the network \mathcal{N}_0 consisting of \mathbf{t} and ∂A and these extra roads has $\text{len}(\mathcal{N}_0) - \text{len}(\mathbf{t}) \leq \frac{3}{2}\text{len}(\partial A)$. It is easy to check that this network \mathcal{N}_0 (without using the edges of \mathbf{t}) satisfies

$$\max_{y \neq y' \in \partial A} \frac{\text{route-length from } y \text{ to } y'}{d(y, y')} \leq 2.$$

In particular, the $K = 0$ case of Lemma 3.3 holds with $\rho(0) = 2$.

Now consider the case $K \geq 1$. To cover the possibility that $\hat{\mathbf{t}}$ and hence \mathbf{t} is entirely in the interior of one of the subrectangles of A , add to \mathcal{N}_0 a road to the boundary from the city closest to the boundary. This road has length at most $\frac{1}{2}a_1 \leq \frac{1}{8}\text{len}(\partial A)$, and the resulting network \mathcal{N}_1 has length $\text{len}(\mathcal{N}_1) - \text{len}(\mathbf{t}) \leq \frac{13}{8}\text{len}(\partial A)$.

Now set

$$\eta := \frac{\frac{13}{8}\text{len}(\partial A)}{K + \binom{K}{2}}.$$

Let \mathcal{N} be the network \mathcal{N}_1 augmented as follows: for each city within distance η from the boundary, add a road from the city to the closest boundary point; for each pair of cities within distance η of each other, add a road directly linking

them. From the definition of η , the extra length added in this stage is at most $\frac{3}{8} \text{len}(\partial A)$, and so \mathcal{N} satisfies the length requirement

$$\text{len}(\mathcal{N}) - \text{len}(\mathbf{t}) \leq 2 \text{len}(\partial A)$$

in Lemma 3.3.

It remain to bound $\text{stretch}^*(\mathcal{N})$. We quote a simple bound on Steiner tree length (given for squares in [2] Lemma 10; the extension to rectangles is straightforward).

Lemma 3.4. *Under the assumptions of Lemma 3.3,*

$$\text{len}(\mathbf{t}) \leq C_1(K) \text{len}(\partial A)$$

where $C_1(K)$ depends only on K .

So $\text{len}(\mathcal{N}) \leq (2 + C_1(K)) \text{len}(\partial A)$ and then

$$\frac{\text{len}(\mathcal{N})}{\eta} \leq \frac{8}{3}(2 + C_1(K))(K + \binom{K}{2}). \quad (16)$$

To bound $\text{stretch}^*(\mathcal{N})$ we need to treat several cases for the pairs (y, y') in (13). We have already obtained an upper bound of 2 for the case where both points are on the boundary. If the two points are at distance $\geq \eta$ apart then, because route length $r(y, y')$ is at most network length, $\frac{r(y, y')}{d(y, y')} \leq \text{len}(\mathcal{N})/\eta$. If the two points are cities within distance η then $\frac{r(y, y')}{d(y, y')} = 1$. The only remaining case is a city y within distance η from the boundary, and a boundary point y' within distance η from y . In this case, by using the edge from y to the closest boundary point and then following the boundary we find (the worst case is near a corner) $\frac{r(y, y')}{d(y, y')} \leq 3$. So

$$\text{stretch}^*(\mathcal{N}) \leq \max(3, \frac{\text{len}(\mathcal{N})}{\eta})$$

and by (16) we have proved Lemma 3.3.

4 Upper bounds on Ψ^{ave} from worst-case stretch for *theta*-graphs

In this section we show how to derive upper bounds on Ψ^{ave} from known bounds on worst-case stretch for graph networks.

The θ_m -graph on a configuration is defined as follows. At each point z_i , consider the natural partition of the plane into m equal-angle cones of base angle $\theta_m = 2\pi/m$ based at z_i ; the boundary lines make angles $(2\pi i/m, 0 \leq i < m)$ with the x -axis. Given z_i and such a cone, each point z_j in the cone has an orthogonal projection onto the bisector line of the cone, at position z'_j say; create an edge (z_i, z_j) for the point z_j in the cone such that z'_j is closest to z_i . This is




Figure 2: Parametrization of z as (r, ℓ) .

a now well-known construction of graphs with low stretch, and the best known explicit bounds on stretch are given as follows in [6]. These hold for $m \geq 6$.

$$\begin{aligned}
1 + \frac{2 \sin(\theta_m/2)}{\cos(\theta_m/2) - \sin(\theta_m/2)} & \quad m = 0 \bmod 4 \\
\frac{\cos(\theta_m/4)}{\cos(\theta_m/2) - \sin(3\theta_m/4)} & \quad m = 1 \text{ or } 3 \bmod 4 \\
1 + 2 \sin(\theta_m/2) & \quad m = 2 \bmod 4.
\end{aligned}$$

Writing s_m for these bounds, we immediately have from (10)

$$\Psi^{\text{ave}}(s_m) \leq L_m \quad (17)$$

where L_m is the mean length-per-unit-area of the θ_m graph over the rate-1 Poisson point process on \mathbb{R}^2 . Calculating L_m is in principle straightforward; indeed more detailed calculations of various statistics in the finite- n random model can be found in [12], though they do not explicitly consider the statistic L_m . We use some of the notation from [12]. As observed there, the calculation is in practice easier in the case of even m , so we treat that case.

Take the origin $\mathbf{0}$ as a typical point of the Poisson configuration. An edge $(\mathbf{0}, z)$ created by the defining rule applied at $\mathbf{0}$ may or may not be *mutual*, meaning it is also created by the rule applied at z . We readily see the formula

$$L_m = m \int_C \|z\| \left(\frac{1}{2} p_{\text{mut}}(z) + p_{\text{not}}(z) \right) dz. \quad (18)$$

Here

C is a cone of base angle $\theta_m = 2\pi/m$;

$\|z\|$ is Euclidean distance from z to the origin;

$p_{\text{mut}}(z)$ is the probability that (if there is a Poisson point at z) there is a mutual edge $(\mathbf{0}, z)$; $p_{\text{not}}(z)$ is the corresponding probability of a non-mutual edge created by the rule at $\mathbf{0}$. The $\frac{1}{2}$ term avoids double-counting mutual edges.

Following [12] we first parameterize a point $z \in C$ by a pair $(r, \ell) = (r(z), \ell(z))$ where, drawing the bisector horizontally in Fig. 2, r and $\ell - r$ are the vertical distances from z to the cone boundaries. If there is a Poisson point at z then the rule at $\mathbf{0}$ creates an edge iff triangle $T = T(z)$ is empty of other Poisson points. Moreover this edge will be mutual iff a certain other triangle $T' = T'(z)$ is also empty. Now [12]

$$\text{area}(T) = \alpha \ell^2; \quad \text{area}(T' \setminus T) = \alpha(r^2 + (\ell - r)^2)$$

where

$$\alpha := \frac{\cos(\theta_m/2)}{4 \sin(\theta_m/2)}.$$

So

$$p_{\text{mut}}(z) = \exp(-\text{area}(T \cup T')) = \exp(-\alpha(\ell^2 + r^2 + (\ell - r)^2))$$

$$p_{\text{not}}(z) = \exp(-\text{area}(T)) - \exp(-\text{area}(T \cup T')) = \exp(-\alpha \ell^2) - \exp(-\alpha(\ell^2 + r^2 + (\ell - r)^2))$$

and so

$$\frac{1}{2}p_{\text{mut}}(z) + p_{\text{not}}(z) = \exp(-\alpha \ell^2) - \frac{1}{2} \exp(-\alpha(\ell^2 + r^2 + (\ell - r)^2)).$$

Substituting into (18) we now have an expression for L_m , for even $m \geq 6$. It is straightforward to show the asymptotics

$$L_m = \Theta(m^{3/2}) \quad s_m - 1 = \Theta(m^{-1})$$

and then from (17) we have

$$\Psi^{\text{ave}}(s) = O((s - 1)^{-3/2}) \text{ as } s \downarrow 1. \quad (19)$$

We will see in section 6 that we can improve this bound when using Steiner networks instead of graph networks.

5 Upper bounds via a “freeways and access roads” construction

In our “Steiner network” setting we can get explicit bounds on Ψ^{worst} via elementary constructions using parallel “freeways” in different directions, with “access roads” linking cities to nearby freeways. We give details in the simplest setting in section 5.1, and state a more general result in section 5.2.

5.1 Constructions based on a square grid of roads

Proposition 5.1.

$$\Psi^{\text{worst}}(2) \leq 4 \quad (20)$$

$$\Psi^{\text{worst}}\left(\frac{3}{2}\right) \leq 4\sqrt{2} \quad (21)$$

$$\Psi^{\text{worst}}(\sqrt{2}) \leq 4\sqrt{3}. \quad (22)$$

Proof. Fix $0 < t_\infty < \infty$ and choose $t = t(n) \rightarrow t_\infty$ such that $n^{1/2}/t(n)$ is an integer $m = m(n)$. First construct a network of *grid roads* which partition the region $[0, n^{1/2}]^2$ into m^2 squares of side-length t . These grid roads (including the boundary of $[0, n^{1/2}]^2$) have total length

$$n^{1/2} \times 2(m+1) \sim 2n/t_\infty.$$

Next, for each city construct a north-south (N-S) and an east-west (E-W) road through the city and across the square containing the city. These *access roads* have total length $2tn$.

We now study the network \mathcal{N}_n^1 thus constructed. We have already seen that

$$n^{-1} \text{len}(\mathcal{N}_n^1) \rightarrow 2(t_\infty + \frac{1}{t_\infty}) \quad (23)$$

so we need to bound the stretch. Note that in a right angle triangle with side-lengths a, b and $c = \sqrt{a^2 + b^2}$ we have

$$\frac{a+b}{c} \leq \sqrt{2}.$$

Thus to show that a city-pair (i, j) has $\frac{r(i, j)}{d(i, j)} \leq \sqrt{2}$ it is enough to show that (supposing w.l.o.g. that city j is to the south-east of city i) there is a route from i to j using only southward and eastward roads. But, consulting Fig. 2, this is clearly true in the three cases

- (i) the two cities are in the same square (as a and b)
- (ii) the two cities are in different rows and different columns (as a and c).
- (iii) the two cities are in adjacent squares (as a and d).

So it remains to consider the final case

- (iv) the two cities are in squares in the same column (say) separated by some number $k \geq 1$ of squares.

The remainder of the argument rests upon being able to recognize, within case (iv), which city positions (v, w) maximize the ratio $r(v, w)/d(v, w)$. In the context of the square grid, these “worst situations” are intuitively clear, and we will state them without proof. It turns out (see Fig. 3, left diagram) that the worst situation in case (iv) is where $k = 1$, this intervening square contains no cities, and the two cities are (arbitrarily close to) the centers of the north and the south edges of the intervening square (as e and f). In this situation $r(v, w)/d(v, w) = 2$, so this is an upper bound for case (iv). Thus the networks \mathcal{N}_n^1 have $\text{stretch}(\mathcal{N}_n^1) \leq 2$. Consulting (23), we can choose $t_\infty = 1$ so that $\text{len}(\mathcal{N}_n^1) \sim 4n$, establishing (20).

Now consider the networks \mathcal{N}_n^2 (Fig. 3, center diagram) obtained from \mathcal{N}_n^1 (left diagram) by adding, for each square, the N-S and the E-W *interior roads* across the square through the center of the square. Now the case (iv) worst situation is where (as g and h in center diagram) the two cities are arbitrarily close to a quarter of the way along the north and the south edges of the intervening square. In this situation $r(g, h)/d(g, h) = 3/2$, so this is an upper bound for case (iv). That is, $\text{stretch}(\mathcal{N}_n^2) \leq 3/2$. The total extra network length is $2n/t$,




Figure 3: All the grid roads and some of the access roads in \mathcal{N}_n^1 .




Figure 4: Networks with grid roads and interior roads, and the positions maximizing $r(v, w)/d(v, w)$. The access roads in Fig. 2 are present but not shown; they are not helpful for these extremal positions.

so $n^{-1} \text{len}(\mathcal{N}_n^2) \rightarrow 2(t_\infty + \frac{2}{t_\infty})$. Choosing $t_\infty = \sqrt{2}$ gives $n^{-1} \text{len}(\mathcal{N}_n^2) \rightarrow 4\sqrt{2}$ and establishes (21).

Finally consider the networks \mathcal{N}_n^3 obtained from \mathcal{N}_n^1 by adding, for each square, two N-S and two E-W *interior roads* partitioning the square into nine equal subsquares. Here the case (iv) worst situation is where (as e and f in Fig. 3, right diagram) the two cities are arbitrarily close to half of the way along the north and the south edges of the intervening square. In this situation $r(e, f)/d(e, f) = 4/3$, so this is an upper bound for case (iv). But here $4/3$ is less than the bound $\sqrt{2}$ from the other cases. So $\text{stretch}(\mathcal{N}_n^3) \leq \sqrt{2}$. The total extra network length (relative to \mathcal{N}_n^1) is $4n/t$, so $n^{-1} \text{len}(\mathcal{N}_n^3) \rightarrow 2(t_\infty + \frac{3}{t_\infty})$. Choosing $t_\infty = \sqrt{3}$ gives $n^{-1} \text{len}(\mathcal{N}_n^3) \rightarrow 4\sqrt{3}$ and establishes (22).

5.2 A generalization

The constructions above were based on horizontal and vertical lines, distance t apart. One can regard that as the $m = 2$ case of the line pattern with m lines through the origin at angles π/m apart, each duplicated by parallel lines distance t apart. Analogous network constructions based on this line pattern were studied in the Master's thesis [10], where it was shown that, for fixed $1 < s < 2$, the construction gives an s -spanner with total length bounded by the quantity $\Psi^*(s)$ below, which is therefore an upper bound on $\Psi^{\text{worst}}(s)$.

Theorem 5.2. *For $1 < s < 2$ set*

$$\phi_s = \frac{\pi}{2} - \sin^{-1}\left(\frac{1}{s}\right)$$

$$\Psi^*(s) = \frac{2\lceil \frac{\pi}{\phi_s} \rceil \sqrt{(1 + \lceil \frac{1}{s-1} \rceil) \tan \phi_s}}{\sin \phi_s}.$$

Then $\Psi^{\text{worst}}(s) \leq \Psi^(s)$.*

In particular, as $s \downarrow 1$ we have $\phi_s \sim \sqrt{2(s-1)}$ and then $\Psi^*(s) \sim 2^{1/4}\pi(s-1)^{-5/4}$, so

$$\Psi^{\text{worst}}(s) = O((s-1)^{-5/4}) \text{ as } s \downarrow 1. \quad (24)$$

We will not repeat the proof of Theorem 5.2 here.

6 Upper bounds by putting a road in every cone

We first show (Proposition 6.1) that one can achieve a given stretch $s > 1$ by insisting that the network has the property of containing roads from each city within each cone of appropriate base angle θ_s . In section 6.2 we show how, in the random model, it is easy to construct networks with the desired property whose expected length can be calculated; this leads to bounds on $\Psi^{\text{ave}}(s)$, stated in Proposition 6.4. This idea is quite similar to the notions of θ -graph from section 4 and of *Yao graph* [17]. But by using Steiner networks instead of graph




Figure 5: A bad configuration for a θ -graph.

networks we obtain in Proposition 6.4 a bound which (for small $s - 1$, at any rate) improves the bound (19) derived from θ -graphs.

To spotlight the essential difference between graph networks and Steiner networks here, Fig. 5 (copied from Fig. 5 of [5]) illustrates a worst-case configuration for stretch of θ -graphs: a route from w to u must go via v or v' . In our construction, there would be a line from w which meets the line (v, u) somewhere near u . It seems plausible that one can get bounds on $\Psi^{\text{worst}}(s)$ in a similar way, adapting other methods from [13], and perhaps improve on Theorem 5.2, but we have not investigated this question carefully.

6.1 The construction

Given a point z in the plane and angles (relative to x -axis, as usual) ϕ and θ , write $\text{cone}(z, \phi, \phi + \theta)$ for the cone bounded by the two rays from z at angles ϕ and $\phi + \theta \bmod 2\pi$. Fix $0 < \theta < \pi/2$. Consider a *graph* network on a given configuration of cities. Call such a network θ -dense if for each city z and each ϕ , if there exists another city in $\text{cone}(z, \phi, \phi + \theta)$, then there exists a road from z to some city in that cone. One can find analogs of Proposition 6.1 below for finite configurations, but it is simpler (and sufficient for our purposes) to work under the assumption

$$\text{for each city } z \text{ and each } \phi, \text{ the } \text{cone}(z, \phi, \phi + \theta) \text{ contains another city} \quad (25)$$

which of course cannot hold for any finite configuration but does hold for the Poisson process on the infinite plane.

Proposition 6.1. *Consider a locally finite configuration on the plane satisfying (25), and consider a θ -dense graph network on the configuration. Then its stretch (considered as a Steiner network) is at most $\frac{1}{\cos \frac{\theta}{2}}$.*

Proof. Fix two cities, w.l.o.g. at $(0, 0)$ and $(x_0, 0)$, where $x_0 > 0$. We first show that the Proposition can be reduced to the following lemma.

Lemma 6.2. *Under the hypotheses of Proposition 6.1, there exist $-\theta < \phi < 0$ and a route from $(0, 0)$ to $(x_0, 0)$ such that the angle of each segment lies in the range $[\phi, \theta + \phi]$.*

Each segment of the Lemma 6.2 route can be visualized as one edge of a triangle whose two other “virtual edges” are at angles ϕ and $\theta + \phi$. So the length of the route is upper bounded by the length of the path of virtual edges for each such triangle (this is a path in the plane, not a route in the network). The length of this path is the length of the path in the plane which goes from $(0, 0)$ to $(x_0, 0)$ by using a line of angle ϕ followed by a line of angle $\theta + \phi$. The length of this path is maximized (as ϕ varies) when $\phi = -\theta/2$ in which case the length equals $x_0 / \cos(\theta/2)$, establishing the Proposition 6.1 bound on stretch.

Proof of Lemma 6.2. Fix $\phi \in (-\theta, 0)$. Define a *lower route* from the city $(0, 0)$ to some point on the line $\{(x_0, y) : -\infty < y < \infty\}$ via the simple procedure: $v_0 = (0, 0)$, and inductively

from v_i , follow the road to v_{i+1} , where v_{i+1} is chosen so that the angle of the segment (v_i, v_{i+1}) is the lowest possible value in $[\phi, \phi + \theta]$ amongst all roads from v_i .

At each step there is some possible choice, by assumption (25) and the assumption of θ -dense. Stop the route where it crosses the line $\{(x_0, y) : -\infty < y < \infty\}$.

Define the analogous *upper route* using the maximum possible angle at each step. It is easy to check that the upper route lies (weakly) above the lower route. In particular, the routes are stopped at two points $(x_0, y_{\text{lower}}^R)$ and $(x_0, y_{\text{upper}}^R)$ where $y_{\text{lower}}^R \leq y_{\text{upper}}^R$. These are *eastward* routes, but we can also define the analogous *westward* routes, which start at city $(x_0, 0)$ and are stopped at points $(0, y_{\text{lower}}^L)$ and $(0, y_{\text{upper}}^L)$ where $y_{\text{lower}}^L \leq y_{\text{upper}}^L$. The roads in these routes are constrained to have angles in the same interval $[\phi, \phi + \theta]$ as in the eastward routes,

To establish the lemma it is enough to show

one of the eastward routes meets one of the westward routes at some point (26)

because then the route from $(0, 0)$ to $(x_0, 0)$ (switching between eastward and westward routes at the meeting point) satisfies the conclusion of the lemma. Clearly (26) holds in the following cases:

- (i) $0 \in [y_{\text{lower}}^R, y_{\text{upper}}^R]$
- (ii) $0 \in [y_{\text{lower}}^L, y_{\text{upper}}^L]$
- (iii) $y_{\text{upper}}^L < 0$ and $y_{\text{upper}}^R < 0$
- (iv) $y_{\text{lower}}^L < 0$ and $y_{\text{lower}}^R < 0$.

There remain two symmetric cases; w.l.o.g. we take the case

- (v) $y_{\text{upper}}^R < 0$ and $y_{\text{lower}}^L > 0$.

The argument so far uses a fixed ϕ ; now we exploit our freedom to choose ϕ . Rewrite y_{lower}^L as $y^L(\phi)$ and rewrite y_{upper}^R as $y^R(\phi)$. We are working in the case: there exists $\phi_0 \in (-\theta, 0)$ such that $y^L(\phi_0) > 0$ and $y^R(\phi_0) < 0$.

Consider the eastward lower route for a given ϕ . The route has some lowest angle, say $\hat{\phi} \geq \phi$. As ϕ increases, the route does not change (and so $y^R(\phi)$ does not change) until ϕ reaches $\hat{\phi}$, at which stage $y^R(\cdot)$ may change but can only increase.

By considering ϕ arbitrarily close to 0, either there is a road from $(0, 0)$ to

$(x_0, 0)$ (in which case the result is trivial) or else $y^R(\phi) > 0$. It follows that there exists some $\phi^* \in [\phi_0, 0)$ such that $y^R(\phi^*) \leq 0$ but $y^R(\phi^* + \varepsilon) \geq 0$ for all sufficiently small $\varepsilon > 0$. Now consider the two eastward lower routes for ϕ^* and for $\phi^* + \varepsilon$. The westward upper route for ϕ^* must meet one of those eastward routes, so the conclusion of the lemma holds for ϕ^* .

6.2 An upper bound on $\Psi^{\text{ave}}(s)$

As in section 2.1 we work with the Poisson process of cities on the infinite plane. There are several ways one might try to use Proposition 6.1; we will just treat one of the simplest. Fix $k \geq 2$. For $0 \leq i \leq k-1$ define a network \mathcal{N}_i by:

for each city z , create a road as a line segment from z to its closest neighbor city in $\text{cone}(z, i\pi/k, (i+1)\pi/k)$, and a road to its closest neighbor city in $\text{cone}(z, \pi + i\pi/k, \pi + (i+1)\pi/k)$.

Network \mathcal{N}_i has a certain normalized length (mean length per unit area) L_k , which by rotational symmetry of the Poisson point process does not depend on i . A calculation below will show

Lemma 6.3.

$$L_k = \sqrt{2k} - \frac{1}{4}\pi^{1/2} \int_0^{\pi/k} \left[\frac{\pi}{k} - \cos \omega \sin \omega + \frac{\sin^2 \omega}{\tan \pi/k} \right]^{-3/2} d\omega. \quad (27)$$

Construct a network \mathcal{N} as the union of \mathcal{N}_i over $0 \leq i \leq k-1$. Its normalized length equals kL_k . And it is clearly θ -dense for $\theta = 2\pi/k$, so Proposition 6.1 bounds its stretch by $1/\cos(\pi/k)$. In other words, using (10)

Proposition 6.4. For each $k \geq 2$, $\Psi^{\text{ave}}(\frac{1}{\cos \frac{\pi}{k}}) \leq kL_k$.

In particular, (27) shows that $kL_k \leq 2^{1/2}k^{3/2}$, whereas $\frac{1}{\cos \frac{\pi}{k}} - 1 \sim \frac{\pi^2}{2k^2}$ as $k \rightarrow \infty$, implying

Corollary 6.5. $\Psi^{\text{ave}}(s) \leq (2^{-1/4}\pi^{3/2} + o(1))(s-1)^{-3/4}$ as $s \downarrow 1$.

Proof of Lemma 6.3. Write $\mathbf{0}$ for the origin. Consider a position measured in polar coordinates as (r, ω) with $0 < \omega < \pi/k$. So $(r, \omega) \in \text{cone}(\mathbf{0}, 0, \pi/k)$ and $\mathbf{0} \in \text{cone}((r, \omega), \pi, \pi + \pi/k)$. Suppose there are cities at $\mathbf{0}$ and at (r, ω) , with other cities at the points of a Poisson process. Define

$$\begin{aligned} p(r, \omega) &= \mathbb{P}((r, \omega) \text{ is nearest city to } \mathbf{0} \text{ in } \text{cone}(\mathbf{0}, 0, \pi/k)) \\ p_1(r, \omega) &= \mathbb{P}((r, \omega) \text{ is nearest city to } \mathbf{0} \text{ in } \text{cone}(\mathbf{0}, 0, \pi/k) \\ &\quad \text{and } \mathbf{0} \text{ is nearest city to } (r, \omega) \text{ in } \text{cone}((r, \omega), \pi, \pi + \pi/k)). \end{aligned}$$

We assert

$$L_k = \int_0^\infty \int_0^{\pi/k} r [2p(r, \omega) - p_1(r, \omega)] d\omega r dr. \quad (28)$$

To argue this, first consider only roads (v_L, v_R) , written so that the x -coordinate of v_L is less than the x -coordinate of v_R , and for which each city is the closest neighbor of the other city in the relevant cone. Given such cities at $v_L = \mathbf{0}$, $v_R = (r, \omega)$ the probability of such a road is $p_1(r, \omega)$ and the contribution to mean network length is $rp_1(r, \omega)$. Because the density of possible positions of (v_L, v_R) has intensity 1 on the region where $v_R \in \text{cone}(v_L, 0, \pi/k)$, the contribution to normalized network length will be

$$\int_0^\infty \int_0^{\pi/k} r [p_1(r, \omega)] d\omega r dr.$$

If instead we consider only roads (v_L, v_R) where v_R is the nearest neighbor to v_L in its cone but not conversely, then similarly the normalized length of such roads is

$$\int_0^\infty \int_0^{\pi/k} r [p(r, \omega) - p_1(r, \omega)] d\omega r dr.$$

By symmetry, the opposite possibility – that v_L is the nearest neighbor to v_R in its cone but not conversely – makes the same contribution. Summing these three contributions gives (28).

To write formulas for $p(\cdot)$ and $p_1(\cdot)$, recall that the probability that the Poisson process assigns no cities to a region A equals $\exp(-\text{area}(A))$. For $p(\cdot)$, the relevant region is the finite cone $\mathbf{0}CE$ in Fig. 6, which has area $\frac{\pi r^2}{2k}$, and so

$$p(r, \omega) = \exp(-\frac{\pi r^2}{2k}). \quad (29)$$

For $p_1(\cdot)$, the relevant region is the entire region $\mathbf{0}ABCDEFGF$ in Fig. 6. The area of this region can be represented as

area of cone $\mathbf{0}CE$, plus area of cone DGA , minus area of parallelogram $\mathbf{0}BDF$.

The parallelogram has height $r \sin \omega$ and base $r \cos \omega - r \frac{\sin \omega}{\tan \pi/k}$ and hence has area

$$r^2 \left(\cos \omega - \frac{\sin \omega}{\tan \pi/k} \right) \sin \omega.$$

So the area of $\mathbf{0}ABCDEFGF$ equals

$$\frac{\pi r^2}{2k} + \frac{\pi r^2}{2k} - r^2 \left(\cos \omega - \frac{\sin \omega}{\tan \pi/k} \right) \sin \omega$$

and finally

$$p_1(r, \omega) = \exp(-r^2 [\frac{\pi}{k} - \cos \omega \sin \omega + \frac{\sin^2 \omega}{\tan \pi/k}]) \quad (30)$$

Returning to formula (28), because $\int_0^\infty r^2 \exp(-ar^2) dr = \frac{1}{4} \pi^{1/2} a^{-3/2}$, we can integrate out r to get

$$\begin{aligned} L_k &= \frac{1}{4} \pi^{1/2} \int_0^{\pi/k} \left(2 \left(\frac{\pi}{2k} \right)^{-3/2} - \left[\frac{\pi}{k} - \cos \omega \sin \omega + \frac{\sin^2 \omega}{\tan \pi/k} \right]^{-3/2} \right) d\omega \\ &= \sqrt{2k} - \frac{1}{4} \pi^{1/2} \int_0^{\pi/k} \left[\frac{\pi}{k} - \cos \omega \sin \omega + \frac{\sin^2 \omega}{\tan \pi/k} \right]^{-3/2} d\omega \end{aligned} \quad (31)$$

which is formula (27).




Figure 6: Regions of integration in the proof of Lemma 6.3.

7 Lower bounds in long networks; average-case analysis

Turning to lower bounds, for Ψ^{ave} we start by giving a reformulation (33) of the interpretation (10) in terms of a Poisson point process on the infinite plane. In (10) we required the distribution μ of the network to be translation invariant; by applying a random rotation Θ (uniform on $(0, 2\pi)$) we may suppose also that μ is isotropic. Recall $L(\mu)$ and $S(\mu)$ denote normalized length and stretch. Consider the number

$$\text{intersect}(\mu) = \text{mean number of intersections of network edges with} \\ \text{the } x\text{-axis per unit length.}$$

There is a general formula (see [16] Chapter 8 for the relevant theory) that for any isotropic translation invariant network,

$$L(\mu) = \frac{\pi}{2} \times \text{intersect}(\mu). \quad (32)$$

So we can rewrite (10) as

$$\Psi^{\text{ave}}(s) = \frac{\pi}{2} \times \inf\{L(\mu); \mu \text{ is isotropic translation invariant, } S(\mu) \leq s\}. \quad (33)$$

We will use this formulation to obtain an order of magnitude lower bound for small s . This general method was used in somewhat different contexts in [2, 1].

Proposition 7.1. $\Psi^{ave}(s) = \Omega((s-1)^{-3/8})$ as $s \downarrow 1$.

Proof. Given $h > 0$ consider the rate-1 Poisson point process restricted to the infinite strip $(-\infty, \infty) \times [-h, h]$. Consider pairs of such Poisson points, where one point is above the x -axis and the other is below the x -axis, and where the line segment between the two points crosses the x -axis at an angle greater than 45° . That is, consider pairs at positions (x_1, y_1) and (x_2, y_2) related by

$$-h < \min(y_1, y_2) < 0 < \max(y_1, y_2) < h; \quad |x_2 - x_1| < |y_2 - y_1|. \quad (34)$$

Call such a pair *friends*. For each friends pair, a hypothetical straight line segment between them crosses the x -axis at some position χ , and the set of all such “virtual crossing positions” is a stationary point process on the line $(-\infty, \infty)$. For $L > 0$ write

$$N(h, L) = \text{number of virtual crossing positions in } [0, L].$$

Now consider a network with stretch $\leq 1 + s$ over the rate-1 Poisson point process on the plane. (So here $s > 0$; this notational shift simplifies formulas). The route between two friends must cross the x -axis at some “route-crossing position” χ' ; write $\delta(h, s)$ for the maximum possible value of the distance between the route-crossing position χ' and the virtual crossing position χ . It is geometrically clear that this maximum is attained when the friends are at positions $(-h, -h)$ and (h, h) , and therefore

$$\delta(h, s) = hg^{-1}(s) \quad (35)$$

where $g^{-1}(\cdot)$ is the inverse function of

$$g(\delta) = \frac{\sqrt{1 + (1 + \delta)^2} + \sqrt{1 + (1 - \delta)^2}}{2\sqrt{2}} - 1$$

for which we calculate

$$g(\delta) \sim \delta^2/8 \text{ as } \delta \downarrow 0. \quad (36)$$

Now choose $L > 0$ and partition the x -axis into blocks of length $L + 2\delta(h, s)$, each block consisting of a middle interval of length L surrounded by two intervals of length $\delta(h, s)$. If the middle interval contains the virtual crossing position for a pair of friends in the Poisson process, then the block contains the route-crossing position, and it follows that the rate of such route-crossing positions is at least $\mathbb{P}(N(h, L) \geq 1)/(L + 2\delta(h, s))$. We may choose h and L arbitrarily, so appealing to (33) we have

$$\Psi^{ave}(1 + s) \geq \frac{\pi}{2} \sup_{h, L} \frac{\mathbb{P}(N(h, L) \geq 1)}{L + 2\delta(h, s)}. \quad (37)$$

We can lower bound the numerator via the second moment inequality

$$\mathbb{P}(N(h, L) \geq 1) \geq \frac{(\mathbb{E}N(h, L))^2}{\mathbb{E}N^2(h, L)}. \quad (38)$$

It is easy to calculate $\mathbb{E}N(h, L)$, as follows. For a point $(0, -y_0)$ consider the set of possible positions of a friend (x, y) with $x > 0$. The constraints are

$$0 < y < h, \quad 0 < x < y_0 + y$$

and the area of this region equals $hy_0 + h^2/2$. It follows easily that the rate of the stationary process of virtual crossing positions equals

$$2 \int_0^h (hy_0 + h^2/2) dy_0 = 2h^3.$$

The initial factor 2 arises due to the symmetric possibility $(0, +y_0)$ for the left point. So we have shown

$$\mathbb{E}N(h, L) = 2h^3L.$$

We will be concerned with the limit regime

$$h \rightarrow \infty, \quad L \rightarrow 0, \quad 2h^3L \rightarrow \lambda \quad (39)$$

for arbitrary $0 < \lambda < \infty$. Intuitively we expect that the distribution of $N(h, L)$ converges to $\text{Poisson}(\lambda)$ in this regime, but for our purposes it will suffice to prove the second moment result (consistent with the Poisson limit)

$$\mathbb{E}N^2(h, L) \rightarrow \lambda^2 + \lambda \text{ in the limit regime (39).} \quad (40)$$

Deferring the proof of (40), Proposition 7.1 can be deduced from the ingredients above. Set

$$h = h(s) = s^{-1/8}, \quad L = L(s) = s^{3/8}$$

and consider orders of magnitude as $s \downarrow 0$. The numerator in (37) is $\Omega(1)$ by (38) and (40). And by (35) and (36) we see that $\delta(h, s)$ is order $hs^{1/2} = s^{3/8}$, so the denominator in (37) is order $s^{3/8}$, establishing the Proposition.

Proof of (40). The formula for the second moment is given as (41) below. The term $\mathbb{E}N(h, L)$ arises from individual crossings, and the term $(\mathbb{E}N(h, L))^2$ is the contribution from pairs of virtual crossing positions in $[0, L]$ for which the 4 end-points are all distinct. The integral term is the contribution from the case of two virtual crossing positions in $[0, L]$ with an end-point in common, say at $(x_0, -y_0)$ where $y_0 > 0$. This term involves the region $A(x_0, -y_0)$ containing the possible positions of a friend of $(x_0, -y_0)$ for which the virtual crossing position is in $[0, L]$. Fig. 7 shows this region, for a particular value of $(x_0, -y_0)$. The integrand $\frac{1}{2}(\text{area } A(x_0, -y_0))^2$ is the mean (conditioned on a point at $(x_0, -y_0)$) number of pairs of friends for which both virtual crossing positions (from friend to $(x_0, -y_0)$) are in $[0, L]$. This leads to the formula

$$\mathbb{E}N^2(h, L) = \mathbb{E}N(h, L) + (\mathbb{E}N(h, L))^2 + 2 \int \int_B \frac{1}{2}(\text{area } A(x_0, -y_0))^2 dx_0 dy_0. \quad (41)$$

We integrate over the region B of values for $(x_0, -y_0)$ which are consistent with a virtual crossing position in $[0, L]$. This region B can be decomposed as the




Figure 7: The region $A(x_0, -y_0)$ for the point \bullet .




Figure 8: The decomposition of the region of points consistent with a virtual crossing position in $[0, L]$. The point \bullet is the same as in Fig. 7.

union of four regions $B_0, B_1^\ell, B_1^r, B_2$ as shown in Fig. 8, wherein we are assuming $h > L/2$, which is true in the limit regime.

For $(x_0, -y_0) \in B_1^r$, the case shown in Fig. 7, the region $A(x_0, -y_0)$ is the trapezoid bounded by the line $y = 0$, the line $y = h$, the line of slope -1 through $(x_0, -y_0)$ and the line through $(x_0, -y_0)$ and $(L, 0)$. A brief calculation shows

$$\text{area } A(x_0, -y_0) = \frac{1}{2} \left(1 + \frac{L-x_0}{y_0} \right) ((h+y_0)^2 - y_0^2) \quad \text{for } (x_0, -y_0) \in B_1^r.$$

Easier calculations show

$$\text{area } A(x_0, -y_0) = (h+y_0)^2 - y_0^2 \quad \text{for } (x_0, -y_0) \in B_0.$$

$$\text{area } A(x_0, -y_0) = \frac{L}{2y_0} ((h+y_0)^2 - y_0^2) \quad \text{for } (x_0, -y_0) \in B_2.$$

The case B_1^ℓ is symmetric with B_1^r . We could calculate $\mathbb{E}N^2(h, L)$ exactly using (41), but we only need an upper bound. The formulas above show that, as x_0 varies for fixed y_0 , the quantity “area $A(x_0, -y_0)$ ” takes its maximum value on B_0 or B_2 , and so

$$\text{area } A(x_0, -y_0) \leq ((h+y_0)^2 - y_0^2) \min(1, \frac{L}{2y_0}).$$

So the integral term in (41) is bounded by

$$\int_0^h (L+2y_0) \left(((h+y_0)^2 - y_0^2) \min(1, \frac{L}{2y_0}) \right)^2 dy_0.$$

The integral over $0 < y_0 < L/2$ works out as $\frac{3}{2}h^4L^2 + \frac{5}{6}h^3L^3 + \frac{7}{24}h^2L^4$. The integral over $L/2 < y_0 < h$ works out as $\frac{7}{2}h^4L^2 - \frac{1}{4}h^3L^3 - \frac{3}{4}h^2L^4 + (\frac{1}{2}L^2h^4 + L^3h^3)\log(2h/L)$. So in the limit regime (39), the leading term is the term $\frac{1}{2}L^2h^4 \log(2h/L)$. But this term $\rightarrow 0$, establishing (40).

8 Lower bounds on Ψ^{worst} based on local optimality

One can get lower bounds on Ψ^{worst} by choosing any configuration of cities and lower bounding the network length required for a network on that particular configuration to have a given stretch. There are heuristic reasons (and the Steiner constant results mentioned at the start of section 3) to suspect that some kinds of regular configurations (rather than typical random configurations) are close to worst-case, so it is not unreasonable to use regular configurations to obtain lower bounds on worst-case behavior. This allows us to work directly on the infinite plane, because the regular configurations we use have known average number of points per unit area.




Figure 9: The “alternate diagonals” network.

8.1 A bound from the square grid

Consider, for instance, the “square grid” configuration of cities at the points $\{(i, j); -\infty < i, j < \infty\}$. The usual “square lattice” network (roads between city pairs (v, w) at distance 1) has normalized length = 2 and stretch = $\sqrt{2}$. It is natural to conjecture this network is optimal amongst Steiner networks, in the following sense.

Conjecture 8.1. *If a Steiner network on the square grid configuration has stretch $\leq \sqrt{2}$ then its normalized length is at least 2.*

If true, this would imply $\Psi^{\text{worst}}(\sqrt{2}) \geq 2$. Similarly, any result of the type

A particular network \mathcal{N}_0 on a particular configuration \mathbf{z} is optimal, in the sense that any other network \mathcal{N} with $\text{stretch}(\mathcal{N}) \leq \text{stretch}(\mathcal{N}_0) = s_0$ has normalized length $L(\mathcal{N}) \geq L(\mathcal{N}_0) = \ell_0$

would imply $\Psi^{\text{worst}}(s_0) \geq \ell_0$. However, we are unable to prove *any* result of this type. Instead, we can only prove weaker results of the following type. Consider the “alternate diagonals” network on the square grid, shown in Fig. 9.

By inspection, this network has normalized length = $\sqrt{2}$ and satisfies

$$\text{route-length from } v \text{ to } w \text{ is } \leq \sqrt{2} \text{ for each city pair } (v, w) \text{ at Euclidean distance 1.} \quad (42)$$

We can prove this network is optimal with respect to those properties.

Proposition 8.2. *Any network on the square grid configuration satisfying (42) has normalized length $\geq \sqrt{2}$.*

Corollary 8.3. $\Psi^{\text{worst}}(\sqrt{2}) \geq \sqrt{2}$.

We call Proposition 8.2 a “local optimality” result because (42) is a “local” analog of stretch.

Proof of Proposition 8.2. Take some network connecting the cities in the square grid configuration. Consider a route through cities $\dots \rightarrow (-2, 0) \rightarrow$

$(-1, 0) \rightarrow (0, 0) \rightarrow (1, 0) \rightarrow (2, 0) \rightarrow \dots$ using minimum-length routes between each successive pair of cities. As we traverse this route, we might *backtrack*, meaning that the x -coordinate of position might decrease, but discarding any backtracking segments leaves a (maybe disconnected) non-backtracking route $((x, y(x)), -\infty < x < \infty)$. Call this “horizontal” route H_0 . Define the measure U_{H_0} on H_0 as the measure induced by Lebesgue measure on x ; that is, a line segment in H_0 from $(x_1, y(x_1))$ to $(x_2, y(x_2))$ has measure $x_2 - x_1$. Repeat for routes H_j through $\dots \rightarrow (-2, j) \rightarrow (-1, j) \rightarrow (0, j) \rightarrow (1, j) \rightarrow (2, j) \rightarrow \dots$. The key observation is that assumption (42) implies that routes H_j are disjoint as j varies, except that routes H_j and H_{j+1} can meet at isolated points of the form $(i + \frac{1}{2}, j + \frac{1}{2})$, as happens in the “alternate diagonals” network.

Let $\mu_H = \sum_{j=-\infty}^{\infty} \mu_{H_j}$. It is clear that μ_H has “density 1”, in the sense that for increasing squares A

$$\frac{\mu_H(A)}{\text{area}(A)} \rightarrow 1 \text{ as } \text{area}(A) \rightarrow \infty. \quad (43)$$

Repeat the construction with vertical routes V_i through $\dots \rightarrow (i, -2) \rightarrow (i, -1) \rightarrow (i, 0) \rightarrow (i, 1) \rightarrow (i, 2) \rightarrow \dots$ to define a measure μ_V which also satisfies (43).

Now write Λ for length measure on the edges of the network. Consider a point (x, y) on a road segment at angle θ . By the disjointness property, this point is in at most one H_j , in which case the density $d\mu_H/d\Lambda$ at the point equals $|\cos \theta|$, and in at most one V_i , in which case the density $d\mu_V/d\Lambda$ at the point equals $|\sin \theta|$. It follows that

$$\frac{d(\mu_H + \mu_V)}{d\Lambda}(x, y) \leq |\cos \theta| + |\sin \theta|. \quad (44)$$

But always $|\cos \theta| + |\sin \theta| \leq \sqrt{2}$, so for any region A

$$\mu_H(A) + \mu_V(A) \leq \sqrt{2} \Lambda(A)$$

and then (43) implies

$$\Lambda(A) \geq (\sqrt{2} - o(1)) \text{area}(A) \text{ as } \text{area}(A) \rightarrow \infty.$$

That is, normalized network length is at least $\sqrt{2}$.

8.2 Another bound from hexagons

Here we show how the argument scheme above can be adapted to the hexagonal configuration of cities (Fig. 11).

Proposition 8.4. *Let \mathcal{N} be a network on the hexagonal configuration such that*

$$\frac{r(v, w)}{d(v, w)} \leq \sqrt{3} \text{ for all (Euclidean) nearest-neighbor pairs } (v, w). \quad (45)$$

Then its normalized length is at least $2^{-1}3^{3/4}$.




Figure 10: An optimality property.

Corollary 8.5. $\Psi^{worst}(\sqrt{3}) \geq 2^{-1}3^{3/4} = 1.14\dots$

Proof of Proposition 8.4. Consider the hexagonal configuration with $\ell =$ distance between nearest neighbors. The density of cities (number per unit area) is

$$\rho(\ell) = 4 \cdot 3^{-3/2} \ell^{-2}. \quad (46)$$

Fig. 10 shows four adjacent cities $ABCD$ in one hexagon. In that figure we see the route lengths satisfy

$$\frac{\text{len}(AZB)}{d(A, B)} = \frac{\text{len}(DZC)}{d(D, C)} = \sqrt{3}$$

and it is easy to check the optimality property:

if π_1 and π_2 are paths in the plane from A to B and from C to D respectively, and if $\max(\frac{\text{len}(\pi_1)}{d(A, B)}, \frac{\text{len}(\pi_2)}{d(D, C)}) \leq \sqrt{3}$, then the paths cannot meet except possibly at Z .

Now consider the minimum-length route in \mathcal{N} through an “angle $\pi/6$ staircase” like $abcdef\dots$ in Fig. 11. By assumption (49) and the optimality property above, this route does not meet the corresponding route through the next staircase $ghijkl\dots$ except at isolated points. As in the previous section, each path segment on such a route is at some angle θ to the “angle = $\pi/6$ ” line; put a measure on the non-backtracking parts of the route with density $\cos \theta$ w.r.t. length measure Λ on the segment. Repeating for each angle = $\pi/6$ staircase gives a measure $\mu_{\pi/6}$ on network edges, which has the property (for squares A)

$$\frac{\mu_{\pi/6}(A)}{\lambda_{\pi/6}(A)} \rightarrow 1 \text{ as } \text{area}(A) \rightarrow \infty$$

where $\lambda_{\pi/6}$ is length measure on the parallel “angle = $\pi/6$ ” *straight* lines through the staircases. The orthogonal distance between such lines equals $3\ell/2$ (this is easiest to see with the angle = $\pi/2$ lines, where the distance is the average of $d(c, d)$ and $d(b, k)$), so

$$\frac{\lambda_{\pi/6}(A)}{\text{area}(A)} \rightarrow \frac{2}{3\ell} \text{ as } \text{area}(A) \rightarrow \infty$$




Figure 11: The hexagonal configuration. Points $abcdef$ are on a “angle = $\pi/6$ staircase” parallel (in an asymptotic sense) to the “angle = $\pi/6$ line”.

and thus

$$\frac{\mu_{\pi/6}(A)}{\text{area}(A)} \rightarrow \frac{2}{3\ell} \text{ as } \text{area}(A) \rightarrow \infty. \quad (47)$$

Repeat the construction with staircases like $cdkl \dots$ with angle = $-\pi/6$ to get a measure $\mu_{-\pi/6}$ on the associated routes; repeat again with staircases like $jkde \dots$ with angle = $-\pi/2$ to get a measure $\mu_{-\pi/2}$. These measures also satisfy (47). Note each adjacent pair of cities is in two staircases, of different angles. The analog of (44) is that, at a point (x, y) on a road segment at angle θ ,

$$\frac{d(\mu_{\pi/2} + \mu_{\pi/6} + \mu_{-\pi/6})}{d\Lambda}(x, y) \leq |\cos(\theta - \pi/2)| + |\cos(\theta - \pi/6)| + |\cos(\theta + \pi/6)|$$

because the point is in at most one route for each of the three angles. But

$$|\cos(\theta - \pi/2)| + |\cos(\theta - \pi/6)| + |\cos(\theta + \pi/6)| \leq 2$$

and so

$$(\mu_{\pi/2} + \mu_{\pi/6} + \mu_{-\pi/6})(A) \leq 2\Lambda(A).$$

Use (47) to see

$$\Lambda(A) \geq (\ell^{-1} - o(1)) \text{area}(A) \text{ as } \text{area}(A) \rightarrow \infty. \quad (48)$$

Our normalization convention is that cities have density 1, that is $\rho(\ell) = 1$ at (46), so $\ell = 2 \cdot 3^{-3/4}$ and the lower bound in (48) becomes $\ell^{-1} = 2^{-1}3^{3/4}$.




Figure 12: The triangular configuration.

8.3 The triangular lattice

We sketch the minor modification which uses the triangular lattice (Fig. 12).

Proposition 8.6. *Let \mathcal{N} be a network on the triangular configuration such that*

$$\frac{r(v, w)}{d(v, w)} \leq \frac{1}{2} + \sqrt{\frac{3}{4}} \text{ for all (Euclidean) nearest-neighbor pairs } (v, w). \quad (49)$$

Then its normalized length is at least $2^{-1/2}3^{3/4}$.

Corollary 8.7. $\Psi^{worst}\left(\sqrt{\frac{3}{4}} + \frac{1}{2}\right) \geq 2^{-1/2}3^{3/4} = 1.61.....$

Outline proof of Proposition 8.6. We indicate changes in the previous argument. The density of cities is now

$$\rho(\ell) = 2 \cdot 3^{-1/2} \ell^{-2}. \quad (50)$$

Fig. 13 shows four adjacent cities $ABCD$ in the triangular configuration. In that figure we see

$$\frac{\text{len}(AZD)}{d(A, D)} = \frac{\text{len}(BZD)}{d(B, C)} = \frac{1}{2} + \sqrt{\frac{3}{4}}$$

and it is easy to check the optimality property:

if π_1 and π_2 are paths in the plane from A to D and from C to B respectively, and if $\max(\frac{\text{len}(\pi_1)}{d(A, D)}, \frac{\text{len}(\pi_2)}{d(B, C)}) \leq \frac{1}{2} + \sqrt{\frac{3}{4}}$, then the paths cannot meet except possibly at Z .




Figure 13: An optimality property.

As before, there is a measure μ_0 on routes through cities like $abckl$ on “angle $= 0$ ” routes, and measures $\mu_{\pi/3}$ and $\mu_{-\pi/3}$ associated with angles $\pi/3$ (like $hicde$) and $-\pi/3$. These satisfy

$$\frac{d(\mu_0 + \mu_{\pi/3} + \mu_{-\pi/3})}{d\Lambda}(x, y) \leq |\cos(\theta)| + |\cos(\theta - \pi/3)| + |\cos(\theta + \pi/3)| \leq 2.$$

The orthogonal distance between parallel lines is $\ell\sqrt{3/4}$, and repeating the argument for (48) leads to

$$\Lambda(A) \geq (3^{1/2}\ell^{-1} - o(1)) \text{ area}(A) \text{ as } \text{area}(A) \rightarrow \infty.$$

Taking $\rho(\ell) = 1$ in (50), the lower bound on normalized length is $3^{1/2}\ell^{-1} = 2^{-1/2}3^{3/4}$.

8.4 Other configurations

One could seek to repeat the arguments above with less symmetric configurations, but the calculations become messier, and we have not pursued details.

Acknowledgments We thank Lisha Li for careful reading of a penultimate draft. Research supported by N.S.F Grant DMS-1106998.

References

- [1] David Aldous. Scale-invariant random spatial networks. *Electron. J. Probab.*, 19:no. 15, 1–41, 2014.
- [2] David J. Aldous and Wilfrid S. Kendall. Short-length routes in low-cost networks via Poisson line patterns. *Adv. in Appl. Probab.*, 40(1):1–21, 2008.
- [3] David J. Aldous and Maxim Krikun. Percolating paths through random points. *ALEA Lat. Am. J. Probab. Math. Stat.*, 1:89–109, 2006.
- [4] David J. Aldous and Julian Shun. Connected spatial networks over random points and a route-length statistic. *Statist. Sci.*, 25(3):275–288, 2010.

- [5] Prosenjit Bose, Jean-Lou De Carufel, Pat Morin, André van Renssen, and Sander Verdonschot. Optimal bounds on theta-graphs: More is not always better. In *Proceedings of the 24th Canadian Conference on Computational Geometry, CCCG 2012, Charlottetown, Prince Edward Island, Canada, August 8-10, 2012*, pages 291–296, 2012.
- [6] Prosenjit Bose, André Renssen, and Sander Verdonschot. On the spanning ratio of theta-graphs. In *Algorithms and Data Structures*, Lecture Notes in Computer Science, pages 182–194. Springer, 2013.
- [7] F. R. K. Chung and R. L. Graham. On Steiner trees for bounded point sets. *Geom. Dedicata*, 11(3):353–361, 1981.
- [8] Vida Dujmovic, Pat Morin, and Michiel H. M. Smid. Average stretch factor: How low does it go?, 2013. <http://arxiv.org/abs/1305.4170>.
- [9] J. Mark Keil and Carl A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, 7(1):13–28, 1992.
- [10] Tamar Lando. Efficient networks and enumerations on forests. Master’s thesis, U.C. Berkeley, 2008. <http://www.stat.berkeley.edu/~aldous/Papers/LandoThesis.pdf>.
- [11] R. E. Miles. On the homogeneous planar Poisson point process. *Math. Biosci.*, 6:85–127, 1970.
- [12] Pat Morin and Sander Verdonschot. On the average number of edges in theta graphs, 2013. <http://arxiv.org/abs/1304.3402>.
- [13] Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, Cambridge, 2007.
- [14] Franco P. Preparata and Michael Ian Shamos. *Computational geometry*. Texts and Monographs in Computer Science. Springer-Verlag, New York, 1985. An introduction.
- [15] J. Michael Steele. *Probability theory and combinatorial optimization*, volume 69 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.
- [16] D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic geometry and its applications*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Ltd., Chichester, 1987. With a foreword by D. G. Kendall.
- [17] A. Yao. On constructing minimum spanning trees in k-dimensional spaces and related problems. *SIAM Journal on Computing*, 11(4):721–736, 1982.

- [18] Joseph E. Yukich. *Probability theory of classical Euclidean optimization problems*, volume 1675 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1998.